

Robust Policy Optimization with Baseline Guarantees

Yinlam Chow*, Marek Petrik†, Mohammad Ghavamzadeh‡

June 17, 2015

Abstract

Our goal is to compute a policy that guarantees improved return over a baseline policy even when the available MDP model is inaccurate. The inaccurate model may be constructed, for example, by system identification techniques when the true model is inaccessible. When the modeling error is large, the standard solution to the constructed model has no performance guarantees with respect to the true model. In this paper we develop algorithms that provide such performance guarantees and show a trade-off between their *complexity* and *conservatism*. Our novel model-based safe policy search algorithms leverage recent advances in robust optimization techniques. Furthermore we illustrate the effectiveness of these algorithms using a numerical example.

1 Introduction

Many problems in science and engineering can be formulated as sequential decision-making under uncertainty. A common scenario in such problems in many different areas, such as online marketing, inventory control, health informatics, and computational finance, is that we are given a batch of data generated by the current strategy(ies) of the company (hospital, investor), and we are asked to find a good or an optimal policy. Although there are many techniques to find a good policy given a batch of data, there are not much results to guarantee that the obtained policy will perform well in the real system without deploying it. Since deploying an untested policy might be risky and harmful for the business, the product (hospital, investment) manager does not allow it unless we provide her with some sort of guarantees on the performance of the policy, e.g., convince her that the policy performs at least as well her existing strategy.

Our focus is on a *model-based* approach to this fundamental problem. In this approach, we first use the batch of data and build a *simulator* that mimics the behavior of the dynamical system under studies (online advertisement, inventory system, emergency room of a hospital, financial market), together with an *error function* that bounds its accuracy, and then use this simulator to generate data and learn a (good) policy. The main challenge here is to have guarantees on the performance of the learned policy,

*Institute of Computational & Mathematical Engineering, Stanford University. Stanford CA, USA 94305

†IBM T. J. Watson Research Center, Yorktown NY, USA 10598

‡Adobe Research, San Jose CA, USA 95110

given the error in the simulator. This line of research is closely related to the area of *robust* learning and control. What makes our problem different than the standard robust learning and control [1, 6, 8, 9, 11, 13, 19, 22] is the existence of a *baseline* policy (e.g., current company’s strategy), which often has a reasonable performance. This difference allows us to develop algorithms, whose performance (the performance of their returned policy) is better than the standard robust methods that optimize for the worst-case scenario.

In this work, we assume that 1) the sequential decision-making problem can be modeled as an *infinite-horizon* Markov decision process (MDP); 2) we are given a *simulator* of this system together with an *error function* that bounds its accuracy (we briefly discuss how the simulator and error function can be built from the batch of data in Appendix A.1); 3) we are given a *baseline policy* for the problem (e.g., the current strategy of the company); and 4) the performance of the baseline policy (*baseline performance*) is known (this is a reasonable assumption as the batch of data is often large enough to have an accurate estimate of the performance of its generating strategy), and our goal is to find a policy that is *safe*, i.e., performs at least as well as the base policy in the real-world. We present four algorithms to tackle this problem in Sections 3.1 to 3.4. For each algorithm, we prove that its returned policy is safe and provide a bound on its performance loss w.r.t. an optimal policy of the real system. From each proposed algorithm to the next, the *computational complexity* grows, but at the same time, the chance of finding a safe policy other than the obvious solution of the baseline policy also increases (the safe policy search become less *conservative*). These major findings are summarized in Figure 1, whose the notations are clearly defined in latter sections. We show this change in the behavior of our algorithms through a simple example in Section 4. This example also serves as a proof of concept for the algorithms in Section 3.1 to 3.4. Another important difference between the algorithms in Section 3.1 to 3.3 and the one in Section 3.4 is that the latter directly works with the baseline policy, while the former uses it in an indirect way, and works with the baseline performance.




Figure 1: An overview of various methods for finding *safe* policies. Here A/R stands for acceptance/rejection of the offline test. The arrow illustrates the computational complexity/conservatism trade-off of different methods.

2 Preliminaries

A γ -discounted MDP is a tuple $\mathcal{M} = \langle \mathcal{X}, \mathcal{A}, r, P, p_0, \gamma \rangle$, where \mathcal{X} and \mathcal{A} are the state and action spaces; $r(x, a) \in [-R_{\max}, R_{\max}]$ is the bounded reward function; $P(\cdot|x, a)$ is the transition probability function; and $p_0(\cdot)$ is the initial state distribution. A solution to MDP \mathcal{M} is a stationary policy π that is a mapping from \mathcal{M} 's state space to a distribution over its action space, i.e., $\pi : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$. We denote by Π_S , the set of all such policies for MDP \mathcal{M} . We define the performance of π in the world modeled by MDP \mathcal{M} as

$$\rho(\pi, \mathcal{M}) = \lim_{T \rightarrow \infty} \mathbb{E} \left(\sum_{t=0}^{T-1} \gamma^t r(X_t, \pi(X_t)) \mid p_0 \right) = p_0^\top V_{\mathcal{M}}^\pi,$$

where X_t is the random variable representing the state of the MDP \mathcal{M} at time-step t and $V_{\mathcal{M}}^\pi$ is the value function of policy π in \mathcal{M} . We also define an optimal policy as $\pi^* \in \arg \max_{\pi \in \Pi_S} \rho(\pi, \mathcal{M})$.

As mentioned in Section 1, we use the historical data to build a *simulator* of the system, together with an error function that measures its accuracy. We denote by \mathcal{M}^* and $\widehat{\mathcal{M}}$ the (unknown) *true* and *simulated* MDPs with transition probability functions P^* and \widehat{P} , respectively.¹ In order to capture the deviation between these models, we make the following assumption:

Assumption 1. *For each $(x, a) \in \mathcal{X} \times \mathcal{A}$, the error function $e(x, a)$ bounds the L_1 -norm of the difference between the true and estimated transition probabilities, i.e.,*

$$\|P^*(\cdot|x, a) - \widehat{P}(\cdot|x, a)\|_1 \leq e(x, a).$$

In many practical situations, the deviation between P^* and \widehat{P} is bounded with high probability. Nevertheless, here we restrict the error bound to be deterministic to simplify the analysis in latter sections. Extending these results with probabilistic bounds is direct and omitted for brevity. More information on building the simulator and computing the L_1 -deviation bound (using an empirical distribution for simulator \widehat{P} and the Weissman distribution bound [21]) is available in Appendix A.1.

Remark 1. *Using the estimated transition probability function \widehat{P} and the error function e , we may construct the uncertainty set*

$$\mathcal{U}(\widehat{P}, e) = \left\{ P : \|P(\cdot|x, a) - \widehat{P}(\cdot|x, a)\|_1 \leq e(x, a), \forall x, a \in \mathcal{X} \times \mathcal{A} \right\}.$$

This uncertainty set automatically defines an uncertainty set $\mathcal{U}(\widehat{\mathcal{M}}, e)$ for MDPs. It is clear that the true transition probability function P^ (the true MDP \mathcal{M}^*) belongs to the uncertainty set $\mathcal{U}(\widehat{P}, e)$.*

Given Assumption 1, we have the following upper-bound on the difference between the return of a policy π in the true and simulated MDPs (with transition probabilities P^* and \widehat{P} , respectively).

¹In this paper, we restrict our attention to error in the transition probability function to simplify the exposition; the results readily extend to the case with error in the reward function.

Lemma 2. *Given Assumption 1, for any policy π , the difference between the return of a policy π in the true and simulated MDPs is upper bounded as follows:*

$$|\rho(\pi, \widehat{\mathcal{M}}) - \rho(\pi, \mathcal{M}^*)| \leq \frac{\gamma R_{\max}}{1 - \gamma} p_0^\top (\mathbf{I} - \gamma P_\pi^*)^{-1} e_\pi,$$

where P_π^* and e_π are the transition probability of the true MDP and the error function when the actions are taken according to policy π . (Appendix A.2)

As discussed in Section 1, we assume that we are provided with a *baseline policy* π_B and we have a very good approximation of its performance $\rho(\pi_B, \mathcal{M}^*)$. We call a policy π *safe* if it is guaranteed to perform not worse than the baseline policy in the true MDP \mathcal{M}^* , i.e., $\rho(\pi, \mathcal{M}^*) \geq \rho(\pi_B, \mathcal{M}^*)$. In the next section, we aim to explore several methods to find a safe policy π , given the simulator $\widehat{\mathcal{M}}$, the error function e , the baseline policy π_B , and the baseline performance $\rho(\pi_B, \mathcal{M}^*)$. For each method, we provide a bound on the performance loss of its returned policy π w.r.t. an optimal policy of the true MDP $\pi_{\mathcal{M}^*}^*$, i.e., $\Phi(\pi) \triangleq \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi, \mathcal{M}^*)$.

3 Computing Safe Policies

In this section, we present several different solutions to finding safe policies in MDP problems, organized from simple to more complex, and discuss their relative advantages and disadvantages.

Before presenting any solution to this problem, let us look at the naive approach of solving the simulated MDP $\widehat{\mathcal{M}}$. Let π_S be an optimal policy of $\widehat{\mathcal{M}}$, i.e., $\pi_S \in \arg \max_{\pi \in \Pi_S} \rho(\pi, \widehat{\mathcal{M}})$ (or equivalently $\pi_S = \pi_{\widehat{\mathcal{M}}}^*$). The following theorem quantifies the performance loss of this policy.

Theorem 3. *Let π_S be an optimal policy of the simulator $\widehat{\mathcal{M}}$. Then under Assumption 1, we have*

$$\Phi(\pi_S) \leq \frac{2\gamma R_{\max}}{(1 - \gamma)^2} \|e\|_\infty. \quad (\text{Appendix A.3})$$

Unfortunately, there is no guarantee that π_S is *safe*, i.e., it performs no worse than the baseline policy π_B . Thus, deploying π_S may lead to undesirable outcomes due to model uncertainties. In the following sections, we present methods whose solutions are guaranteed to be safe.

3.1 Solution based on a Reward Adjusted MDP

In this section, we propose a method that relies on solving the MDP $\widetilde{\mathcal{M}} = \langle \mathcal{X}, \mathcal{A}, \widehat{r}, \widehat{P}, p_0, \gamma \rangle$, which is exactly the same as the simulated MDP $\widehat{\mathcal{M}}$, except that its reward function is *adjusted* as

$$\widehat{r}(x, a) = r(x, a) - \frac{\gamma R_{\max}}{1 - \gamma} e(x, a), \quad \forall x \in \mathcal{X}, \forall a \in \mathcal{A}. \quad (1)$$

The unique property of this MDP is that, when Assumption 1 holds, the performance of any policy π in $\widetilde{\mathcal{M}}$ is a lower-bound on its performance in the true MDP \mathcal{M}^* ,

i.e., $\rho(\pi, \widetilde{\mathcal{M}}) \leq \rho(\pi, \mathcal{M}^*)$ (see Theorem 4). Algorithm 1 summarizes the method for computing a policy using the reward adjusted MDP (RaMDP) $\widetilde{\mathcal{M}}$. It returns an optimal policy of $\widetilde{\mathcal{M}}$, when the performance of this policy in $\widetilde{\mathcal{M}}$ is better than the baseline performance $\rho(\pi_B, \mathcal{M}^*)$, and returns π_B otherwise.

Algorithm 1: Solution based on the RaMDP

input : Simulated MDP $\widetilde{\mathcal{M}}$, baseline performance $\rho(\pi_B, \mathcal{M}^*)$ and error function e

output: Policy π_{Ra}

- 1 $\widehat{r}(x, a) \leftarrow r(x, a) - \frac{\gamma R_{\max}}{1-\gamma} e(x, a)$;
 - 2 $\pi_0 \leftarrow \arg \max_{\pi \in \Pi_S} \rho(\pi, \widetilde{\mathcal{M}})$;
 - 3 $\rho_0 \leftarrow \rho(\pi_0, \widetilde{\mathcal{M}})$;
 - 4 If $\rho_0 > \rho(\pi_B, \mathcal{M}^*)$ $\pi_{Ra} \leftarrow \pi_0$ $\pi_{Ra} \leftarrow \pi_B$ **return** π_{Ra}
-

Since the performance of any policy in $\widetilde{\mathcal{M}}$ is a lower-bound on its performance in \mathcal{M}^* , it guarantes that the solution policy returned by Algorithm 1, π_{Ra} , performs at least as well as the baseline policy π_B . Theorem 4 shows that π_{Ra} is a safe policy and quantifies its performance loss.

Theorem 4. *Given Assumption 1, the solution π_{Ra} of Algorithm 1 is safe, i.e., $\rho(\pi_{Ra}, \mathcal{M}^*) \geq \rho(\pi_B, \mathcal{M}^*)$. Moreover, its performance loss $\Phi(\pi_{Ra})$ satisfies*

$$\Phi(\pi_{Ra}) \leq \min \left\{ \frac{2\gamma R_{\max}}{(1-\gamma)^2} \|e_{\pi_{\mathcal{M}^*}^*}\|_{1, u_{\mathcal{M}^*}^*}, \Phi(\pi_B) \right\},$$

where $u_{\mathcal{M}^*}^*$ is the normalized state occupancy frequency of the optimal policy $\pi_{\mathcal{M}^*}^*$. (Appendix A.4)

Note that Theorem 4 indicates that by this simple adjustment in the reward function of the simulated MDP $\widetilde{\mathcal{M}}$, we may guarantee that our solution is safe. Moreover, it shows that the bound on the performance loss of π_{Ra} is actually tighter than that for the solution π_S of the simulator $\widetilde{\mathcal{M}}$ in Theorem 3. In particular, the L_∞ -norm has been replaced by a weighted L_1 -norm. In terms of computational complexity, since Algorithm 1 only requires solving a standard MDP, it can be implemented by either *value iteration*, *policy iteration* or *linear programming* [18]. The corresponding complexity is therefore $O(|\mathcal{A}||\mathcal{X}|^2/(1-\gamma))$ (for value iteration) [14], which is low-polynomial in $|\mathcal{A}|$, $|\mathcal{X}|$ and $1/(1-\gamma)$.

While Algorithm 1 provides good theoretical guarantees and has low costs of computation, it may be overly conservative in many circumstances. This is because the adjustment of the reward function is based on the assumption that there exists a state with the optimal value of $R_{\max}/(1-\gamma)$ and that this state is accessible from each other state with the reward R_{\max} . Since this assumption is rarely true, we propose a more adaptive formulation in the following section via RMDP methods.

3.2 Solution based on the Robust MDP

Robust optimization is a standard technique to deal with model uncertainty. In this section, we propose an algorithm for finding a safe policy via solving the robust MDP (RMDP). We prove that the policy returned by this algorithm is safe and has better (sharper) worst-case guarantees.

Algorithm 2: Solution based on the RMDP

input : Simulated MDP $\widehat{\mathcal{M}}$, baseline performance $\rho(\pi_B, \mathcal{M}^*)$ and the error function e
output: Policy π_R
 // Construct the uncertainty set
 1 $\mathcal{U}(\widehat{P}, e) = \left\{ P : \|P(\cdot|x, a) - \widehat{P}(\cdot|x, a)\|_1 \leq e(x, a), \forall x, a \in \mathcal{X} \times \mathcal{A} \right\};$
 2 $\pi_0 \leftarrow \arg \max_{\pi \in \Pi_S} \min_{P \in \mathcal{U}(\widehat{P}, e)} \rho(\pi, \mathcal{M}(P));$
 3 $\rho_0 \leftarrow \min_{P \in \mathcal{U}(\widehat{P}, e)} \rho(\pi_0, \mathcal{M}(P));$
 4 If $\rho_0 > \rho(\pi_B, \mathcal{M}^*)$ $\pi_R \leftarrow \pi_0$ $\pi_R \leftarrow \pi_B$ **return** π_R

The robust method is summarized in Algorithm 2. It first constructs an uncertainty set $\mathcal{U}(\widehat{P}, e)$ using the simulator $\widehat{\mathcal{M}}$ and error function e . It then solves the resultant RMDP ($\mathcal{M}(P)$ with $P \in \mathcal{U}(\widehat{P}, e)$) and returns its solution, if its worst-case performance over the uncertainty set is better than the baseline performance $\rho(\pi_B, \mathcal{M}^*)$, and returns π_B otherwise.

Algorithm 2 involves solving a (s, a) -rectangular RMDP. RMDPs satisfy many of the same properties as regular MDPs, such as the existence of an optimal stationary policies. RMDPs can often be solved quite efficiently by value iteration, policy iteration, or modified policy iteration [9, 12, 16, 22]. In practical terms, solving an RMDP can be expensive due to the need to compute the worst-case transition probability which relies on solving a convex optimization problem for each state and action in every iteration. However, the robust solution can be computed very efficiently when the uncertainty set is described, as in our formulation, in terms of an L_1 -norm [17]; similar results also exist for the L_2 norm [9]. Since implementing Algorithm 2 involves solving RMDPs, the complexity is $O(|\mathcal{A}||\mathcal{X}|^3 \log(|\mathcal{X}|)/(1-\gamma))$ (for robust value iteration) [16]. It is higher than that in Algorithm 1 but is still polynomial in $|\mathcal{A}|$, $|\mathcal{X}|$ and $1/(1-\gamma)$. The following theorem shows that the policy π_R is safe and quantifies its performance loss.

Theorem 5. *Given 1, the nonempty solution π_R of Algorithm 2 is safe, i.e., $\rho(\pi_R, \mathcal{M}^*) \geq \rho(\pi_B, \mathcal{M}^*)$. Moreover, its performance loss $\Phi(\pi_R)$ satisfies*

$$\Phi(\pi_R) \leq \min \left\{ \frac{2\gamma R_{\max}}{(1-\gamma)^2} \|e^{\pi_{\mathcal{M}^*}^*}\|_{1, u_{\mathcal{M}^*}^*}, \Phi(\pi_B) \right\},$$

where $u_{\mathcal{M}^*}^*$ is the normalized state occupancy frequency of the optimal policy $\pi_{\mathcal{M}^*}^*$. (Appendix A.5)

Compared to π_S and the bound in Theorem 3 on its performance loss, Theorem 5 indicates that the policy π_R returned by Algorithm 2 is safe and has a smaller bound on

its performance loss. In particular, the bound depends only on a weighted L_1 -norm of the errors for the optimal policy, instead of the L_∞ -norm over all policies.

While the complexity of Algorithm 2 is higher than that of Algorithm 1 (because solving RMDPs is more complex than standard MDP), Theorem 5 does not show any advantage for π_R over the policy π_{RA} returned by Algorithm 1, neither in terms of safety nor in terms of the bound on its performance loss (Theorem 4). This arises the question that why should one use the more complex Algorithm 2 in place of Algorithm 1. Proposition 6 provides an answer to this question and shows whenever Algorithm 1 returns a safe policy, so does Algorithm 2, while the converse is not necessarily true. This implies with extra computational complexity, the conservatism of safe policy search decreases.

Proposition 6. *Given 1, for each policy π , we have*

$$\min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)) \geq \rho(\pi, \widetilde{\mathcal{M}}),$$

where $\widetilde{\mathcal{M}}$ is defined in Section 3.1.

(A.6)

Note that the bound in π_R is based on the worst-case transition probabilities in $\mathcal{U}(\hat{P}, e)$. While this approach guarantees to yield a safe policy, it may still be too conservative as solving a RMDP considers finding a safe policy under the worst-case scenario. Next, we investigate an alternative approach to return a safe but less conservative policy.

3.3 Solution based on an Augmented Robust MDP

As discussed at the end of Section 3.2, the goal in this section is to develop a new method that combines simulated and RMDPs, and reduces the conservatism of *safe policy search* compared to Algorithm 2. We start this section by considering the following constraint optimization problem that finds a policy that maximizes the performance in the simulator and satisfies the safety constraint:

$$\max_{\pi \in \Pi_H} \rho(\pi, \widehat{\mathcal{M}}), \quad \text{subject to} \quad \rho(\pi, \mathcal{M}^*) \geq \rho(\pi_B, \mathcal{M}^*), \quad (2)$$

where Π_H is the general set of history-based policies. To solve (2), we employ the Lagrangian relaxation procedure [3] to convert it to the following unconstrained optimization problem:

$$\max_{\pi \in \Pi_H} \min_{\lambda \geq 0} L(\pi, \lambda) := \rho(\pi, \widehat{\mathcal{M}}) + \lambda \left(\rho(\pi, \mathcal{M}^*) - \rho(\pi_B, \mathcal{M}^*) \right), \quad (3)$$

where λ is the Lagrange multiplier. Unfortunately, solving the optimization problem (2) is impossible, since the true MDP \mathcal{M}^* is unknown. Before describing how we tackle this issue, let us define a few terms and quantities. For the simulated MDP $\widehat{\mathcal{M}}$ and any MDP $\mathcal{M}(P)$ that is only different with $\widehat{\mathcal{M}}$ in its transition probability function, and for any $\lambda_1, \lambda_2 \geq 0$, we define the *augmented* MDP $\mathcal{M}_{\lambda_1, \lambda_2}^A = \langle \mathcal{X} \times \mathcal{X}, \mathcal{A}, r^A, P^A, p_0^A, \gamma \rangle$, where $r^A(x, y, a) = \lambda_1 r(x, a) + \lambda_2 r(y, a)$, $P^A(x', y' | x, y, a) =$

$P(x'|x, a)\hat{P}(y'|y, a)$, and $p_0^A(x, y) = p_0(x)p_0(y)$. Here we can think of x 's and y 's as the states evolved according to the MDPs $\mathcal{M}(P)$ and $\hat{\mathcal{M}}$, respectively. We denote by Π_S^A the set of stationary Markovian policies over this augmented MDP. Using the augmented MDP, for any policy $\pi \in \Pi_H$, we define the Lagrangian function

$$L_P(\pi, \lambda) = \rho(\pi, \mathcal{M}_{\lambda,1}^A(P)) - \lambda \rho(\pi_B, \mathcal{M}^*). \quad (4)$$

Using the above definitions, the Lagrangian function in (3) may be easily written as $L(\pi, \lambda) = L_{P^*}(\pi, \lambda)$. From Theorem 3.6 in [2], it can be easily shown that for any transition probability function $P \in \mathcal{U}(\hat{P}, e)$, the following strong duality holds:

$$\max_{\pi \in \Pi_H} \min_{\lambda \geq 0} L_P(\pi, \lambda) = \max_{\pi \in \Pi_S^A} \min_{\lambda \geq 0} L_P(\pi, \lambda) = \min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} L_P(\pi, \lambda). \quad (5)$$

Setting $P = P^*$, (5) implies that we can replace the class of history-based policies Π_H with the set of stationary Markovian policies over the augmented MDP Π_S^A in problem (3). This further means that Π_S^A is the class of *dominating* policies for the optimization problem (2).

Recall that the duality theory [3] indicates that if the dual Lagrangian problem is bounded, the primal Lagrangian problem is always feasible. This means that if we can find a lower-bound for the dual Lagrangian problem $\min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} L(\pi, \lambda)$ (note that $L(\pi, \lambda) = L_{P^*}(\pi, \lambda)$ and $P^* \in \mathcal{U}(\hat{P}, e)$), the corresponding policy will be feasible for the constraint in (2), which itself means that it is safe. This motivates us to find a saddle-point for the (augmented) robust optimization problem

$$\min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda). \quad (6)$$

Note that compared to the Lagrangian function $L(\pi, \lambda)$, in (6), we have replaced the true transition probability P^* with the worst-case transition probability over the uncertainty set $\mathcal{U}(\hat{P}, e)$. The reason for finding a saddle-point of (6) is because the solution is a lower-bound for the dual Lagrangian:

$$\min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda) \stackrel{(a)}{=} \min_{P \in \mathcal{U}(\hat{P}, e)} \min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} L_P(\pi, \lambda) \stackrel{(b)}{\geq} \min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} L(\pi, \lambda),$$

(a) Theorem 1 in [16] shows that strong duality holds in $(\mathcal{X} \times \mathcal{A})$ -rectangular robust optimization problems, i.e.,

$$\min_{P \in \mathcal{U}(\hat{P}, e)} \max_{\pi \in \Pi_S^A} L_P(\pi, \lambda) = \max_{\pi \in \Pi_S^A} \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda),$$

(b) This is from the fact that $L(\pi, \lambda) = L_{P^*}(\pi, \lambda)$ and $P^* \in \mathcal{U}(\hat{P}, e)$.

Therefore, if we find a saddle point (π_0, λ^*) of the (augmented) robust optimization problem (6), then the corresponding policy π_0 is safe. Given the above observations, we now present Algorithm 3 and prove in Theorem 7 that the policy returned by this algorithm is safe and quantify its performance loss. On Line 2 of Algorithm 3, we use the conventional sub gradient descent approach to solve for a saddle-point. In this approach, we first fix the Lagrange multiplier and solve for an optimal stationary policy

and then optimize for the Lagrangian multiplier (which is a convex optimization problem). These two steps are repeated until the solution converges to a saddle point. More details of this procedure can be found in Appendix A.7. Regarding to the computation cost of this approach, since each sub-gradient descent step involves a RMDP, similar to Algorithm 2 it has complexity of $O(|\mathcal{A}||\mathcal{X}|^3 \log(|\mathcal{X}|)/(1-\gamma))$ for robust value iteration. Thus the total complexity of Algorithm 3 is $O(|\mathcal{A}||\mathcal{X}|^3 \log(|\mathcal{X}|)/(1-\gamma)(1/\sqrt{K}))$ [4], where K is the number of iteration of sub-gradient descent and $O(1/\sqrt{K})$ is the standard convergence rate for first order methods in convex optimization.

Regarding to the augmented Markovian policy π_0 in Algorithm 3, since π_0 requires state information from both the uncertain MDP $\mathcal{M}(P)$ (with state X_t) and simulated MDP $\widehat{\mathcal{M}}$ (with state Y_t). Therefore implementing this policy in real-world (i.e. $P = P^*$) requires real-time state trajectories from the online simulator as well. This inevitably increases the complexity of implementation.

Algorithm 3: Solution based on the Augmented RMDP

input : Simulated MDP $\widehat{\mathcal{M}}$, baseline performance $\rho(\pi_B, \mathcal{M}^*)$ and the error function e
output: Policy π_{AR}
1 Construct the uncertainty set $\mathcal{U}(\widehat{P}, e)$ and augmented MDP $\mathcal{M}_{\lambda,1}^A(P)$, $\forall P \in \mathcal{U}(\widehat{P}, e)$;
2 Solve $\min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} \min_{P \in \mathcal{U}(\widehat{P}, e)} L_P(\pi, \lambda)$ for a saddle-point (π_0, λ^*) ;
3 If a saddle point solution (π_0, λ^*) exists $\pi_{AR} \leftarrow \pi_0$ $\pi_{AR} \leftarrow \pi_B$ **return** π_{AR}

Theorem 7. *Given Assumption 1, the nonempty solution π_{AR} of Algorithm 3 is safe, i.e., $\rho(\pi_{RS}, \mathcal{M}^*) \geq \rho(\pi_B, \mathcal{M}^*)$. Moreover, its performance loss $\Phi(\pi_{AR})$ satisfies*

$$\Phi(\pi_{AR}) \leq \min \left\{ \frac{2\gamma R_{\max}}{(1-\gamma)^2} \|e_{\pi_{\mathcal{M}^*}^*}\|_{1, u_{\mathcal{M}^*}^*}, \Phi(\pi_B) \right\},$$

where $u_{\mathcal{M}^*}^*$ is the normalized state occupancy frequency of the optimal policy $\pi_{\mathcal{M}^*}^*$. (Appendix A.8)

Similar to Section 3.2, compared to π_S and the bound in Theorem 3 on its performance loss, Theorem 7 indicates that the policy π_{AR} returned by Algorithm 3 is safe and has a tighter bound on its performance loss. However, while Algorithm 3 has higher computational complexity than Algorithm 1 and Algorithm 2, Theorem 7 does not show any advantage for its returned policy π_{AR} , over those returned by Algorithm 1 and Algorithm 2, π_{RA} and π_R . This raises a question similar to that in Section 3.2 that why should we use Algorithm 3 instead of Algorithm 2 then? Proposition 8 provides an answer to this question and shows whenever Algorithm 2 returns a safe policy, so does Algorithm 3, while the converse is not always true. This again resonates with the fact that with extra computational complexity, we reduce conservatism of safe policy search.

Proposition 8. *Given Assumption 1, if $\rho(\pi_B, \mathcal{M}^*) \leq \min_{P \in \mathcal{U}(\hat{P})} \rho(\pi_R, \mathcal{M}(P))$, then $\min_{P \in \mathcal{U}(\hat{P})} L_P(\pi_{AR}, \lambda^*)$ is lower-bounded. This means that if Algorithm 2 returns a safe policy other than π_B , so does Algorithm 3. (Appendix A.9)*

3.4 Combining Robust and Baseline Policies

While the robust solution described in Algorithm 2 (or Algorithm 3) is less conservative than Algorithm 1 it may nevertheless be overly restrictive. This is because the proposed improved policy π_0 is evaluated for the worst-case realization of the transition probability in $\mathcal{U}(\hat{P}, e)$ while the return of the baseline policy is with respect to P^* . As a result, a candidate policy π_0 can be rejected even if for any realization $P \in \mathcal{U}(\hat{P}, e)$ the policy is better than the baseline. The left example in Figure 2 depicts the case in which the robust solution will be too restrictive.




Figure 2: Left: Example in which the policy returned by Algorithm 2 is too restrictive. Right: Example in which the policy returned by Algorithm 2 performs better than the baseline policy in some states but not in others.

An additional limitation from Algorithm 1, 2 and 3 is that when the evaluation criterion of the computed policy fails, then the baseline policy is not improved at all. The restriction of these approaches is illustrated by the following counter example:

Consider a simple case when the MDP \mathcal{M} is composed of two separate MDPs \mathcal{M}_1 and \mathcal{M}_2 and suppose that the estimated model of \mathcal{M}_1 is very good, while the model of \mathcal{M}_2 is quite imprecise. Also assume that the initial distribution is uniformly distributed between an initial state in \mathcal{M}_1 and an initial state in \mathcal{M}_2 . Intuitively, the best solution would use the optimized policy for \mathcal{M}_1 , which has a precise model, and use the baseline policy for \mathcal{M}_2 , which has an imprecise model. However, Algorithm 2 simply returns the baseline policy for both components because the return in \mathcal{M}_2 potentially reduces the quality of the robust solution. This phenomenon is illustrated in the right example of Figure 2 for which it would be beneficial to combine the baseline policy with the optimized one instead of returning either one.

The above limitations can be solved by modifying the robust optimization problem. Intuitively, the objective is to find a policy that maximizes improvement over the baseline for any plausible transition probabilities. Equivalently this reduces to a distributionally robust optimization (DRO):

$$\pi_I \in \arg \max_{\pi} \min_{P \in \mathcal{U}(\hat{P}, e)} \left(\rho(\pi, P) - \rho(\pi_B, P) \right). \quad (7)$$

Compared with Algorithm 1, 2 and 3, the first appealing fact of this approach is that the solution policy is always safe because π_B is feasible to (7). The second appealing




Figure 3: Improvement in return over the baseline policy for the proposed methods.

fact is that it only requires the knowledge of baseline policy π_B , without the baseline performance $\rho(\pi_B, P^*)$. Nevertheless, (7) is a DRO problem which in general could be NP-hard to solve [5]. Here the computational methods used to solve this problem are beyond the scope of this paper; we approximate the solution using a heuristic iterative algorithm based on value iteration. The additional complexity of this formulation again corroborates with the phenomenon that additional computational complexity reduces conservatism in safe policy search. Similar to the other three formulations, the following theorem states the safety of the computed policy and describes its performance loss.

Theorem 9. *Given that Assumption 1 is satisfied, then a solution π_I to (7) is safe, i.e., $\rho(\pi_I, P^*) \geq \rho(\pi_B, P^*)$. Moreover, the performance loss of π_I satisfies:*

$$\Phi(\pi_I) \leq \min \left\{ \frac{2\gamma R_{\max}}{(1-\gamma)^2} \left(\|e_{\pi_{\mathcal{M}^*}}\|_{1, u_{\mathcal{M}^*}^*} + \|e_{\pi_B}\|_{1, u_{B, \mathcal{M}^*}} \right), \Phi(\pi_B) \right\},$$

where $u_{\mathcal{M}^*}^*, u_{B, \mathcal{M}^*}$ are the normalized state occupancy frequencies of the optimal and base policies $\pi_{\mathcal{M}^*}^*, \pi_B$ respectively. Also, this bound is tight. (Appendix A.10)

4 Numerical Comparison

In this section, we numerically evaluate the proposed methods on a synthetic benchmark MDP. The benchmark problem loosely models customer interactions with an online system. The four available actions influence the user behavior along two dimensions. Rewards, which represent user satisfactions, vary only along the first dimension. The second dimension influences only transition probabilities. To simulate a realistic source of a baseline policy, we construct it to be optimal when the second dimension of the MDP is ignored. The simulator is constructed directly from the empirical transition probabilities. The transition error e is based on the sampling bounds in Section A.1 and decreases with a square root of the number of samples. Figure 3 depicts the percentage improvement in total return over the baseline policy as a function of the overall number of samples used in constructing the simulator. The methods used in the comparison are as follows. The dashed line shows the return of the optimal policy. EXP stands for

the standard MDP method with the expected return objective. For a small number of samples, this standard method does significantly worse than the baseline policy. RWA stands for the method in Algorithm 1, which leads to a safe policy but, as expected, is overly conservative. ROB stands for the robust method in Algorithm 2, which also guarantees the safety of returned policies, but is much less conservative than RWA. Finally, RBC is the algorithm described in Section 3.4. RBC optimizes the policy in states with many samples and falls back onto the baseline policy otherwise. The combined policy of RBC is not only safe, but also significantly improves on the baseline policy even when the number of samples is small.

5 Conclusion

In this paper we presented four model based safe policy search methods and analyzed their performance. Ranging from computational complexity to conservatism, our approaches provide a full gamut of tools to design good policies offline that match baseline performance. To the best of our knowledge, this line of work is *novel* in the RL community. Similar approaches in the model-free setup can be found in [7, 10, 20], where safe policy evaluation takes place during exploration.

On the technical side, an important future direction is to compare the performance of policies generated by different safe policy search algorithms and to explicitly study the solution algorithm in Section 3.4. On the experimental side, future work includes running advanced simulations in realistic domains such as battery charging/discharging control for smart grid systems.

Bibliography

- [1] A. Ahmed, P. Varakantham, Y. Adulyasak, and P. Jaillet. Regret based robust solutions for uncertain Markov decision processes. In *Advances in Neural Information Processing Systems*, pages 881–889, 2013.
- [2] E. Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- [3] D. Bertsekas. *Nonlinear programming*. Athena Scientific, 1999.
- [4] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [5] E. Delage and Y. Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612, 2010.
- [6] E. Even-Dar, S. Mannor, and Y. Mansour. Action elimination and stopping conditions for reinforcement learning. In *ICML*, pages 162–169, 2003.
- [7] M. Ghavamzadeh and A. Lazaric. Conservative and greedy approaches to classification-based policy iteration. In *AAAI-26th Conference on Artificial Intelligence*, 2012.
- [8] G. Hanasusanto, D. Kuhn, and W. Wiesemann. Two-stage robust integer programming. *Optimization Online*, 2014.
- [9] G. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.

- [10] S. Kakade and J. Langford. Approximately optimal approximate reinforcement learning. In *ICML*, volume 2, pages 267–274, 2002.
- [11] S. Kalyanasundaram, E. Chong, and N. Shroff. Markov decision processes with uncertain transition rates: Sensitivity and robust control. In *Decision and Control, 2002, Proceedings of the 41st IEEE Conference on*, volume 4, pages 3799–3804. IEEE, 2002.
- [12] D. Kaufman and A. Schaefer. Robust modified policy iteration. *INFORMS Journal on Computing*, 25(3):396–410, 2013.
- [13] B. Li and S. Si. Robust dynamic programming for discounted infinite-horizon Markov decision processes with uncertain stationary transition matrices. In *ADPRL*, pages 96–102. IEEE, 2007.
- [14] M. Littman, T. Dean, and L. Kaelbling. On the complexity of solving Markov decision problems. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pages 394–402. Morgan Kaufmann Publishers Inc., 1995.
- [15] P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- [16] A. Nilim and L. El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.
- [17] M. Petrik. RAAM : The benefits of robustness in approximating aggregated MDPs in reinforcement learning. In *Neural Information Processing Systems (NIPS)*, 2014.
- [18] M. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [19] Y. Le Tallec. *Robust, risk-sensitive, and data-driven control of Markov decision processes*. PhD thesis, Massachusetts Institute of Technology, 2007.
- [20] P. Thomas, G. Theodorou, and M. Ghavamzadeh. High confidence off-policy evaluation. *arXiv*, 2015.
- [21] T. Weissman, E. Ordentlich, G. Seroussi, S. Verdu, and M. Weinberger. Inequalities for the L_1 deviation of the empirical distribution. *Hewlett-Packard Labs, Tech. Rep.*, 2003.
- [22] W. Wiesemann, D. Kuhn, and B. Rustem. Robust Markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013.

A Proofs

A.1 Sampling Bounds

[22] showed that L_1 -deviation of the empirical distribution from the true distribution over m distinct events from n samples is bounded as

$$\mathbb{P} \left\{ \|P(\cdot) - \hat{P}(\cdot)\|_1 \geq \epsilon \right\} \leq (2^m - 2) \exp \left(-\frac{n\epsilon^2}{2} \right). \quad (8)$$

Now consider a fixed state-action pair (x, a) and assume that the transition probability $\hat{P}(\cdot|x, a)$ has been estimated using $N(x, a)$ visits to (x, a) . The random event for the transition probability estimate is the state to which the system transits. In this case, $m = |\mathcal{X}|$ and using (8) we may write

$$\|P(\cdot|x, a) - \hat{P}(\cdot|x, a)\|_1 \leq \sqrt{\frac{2}{N(x, a)} \log \left(\frac{2^{|\mathcal{X}|} - 2}{\delta} \right)},$$

with probability at least $1 - \delta$. Thus, by setting $e(x, a) = \sqrt{\frac{2}{N(x, a)} \log \left(\frac{|\mathcal{X}| |\mathcal{A}| (2^{|\mathcal{X}|} - 2)}{\delta'} \right)}$, we can guarantee that $\mathbb{P} \left\{ P^* \notin \mathcal{U}(\hat{P}) \right\} \leq \delta$.

A.2 Proof of Lemma 2

Since the return of a policy π is the product of the initial state distribution and the value function of the policy, i.e., $\rho(\pi, \mathcal{M}) = p_0^\top V_{\mathcal{M}}^\pi$, Lemma 2 is a direct consequence of the following lemma.

Lemma 10. *Consider two MDPs \mathcal{M}_1 and \mathcal{M}_2 that are only different in their transition probability functions P_1 and P_2 , and reward functions r_1 and r_2 . Let π_1 be a policy in \mathcal{M}_1 and π_2 be a policy in \mathcal{M}_2 . Under the assumption that for any state $x \in \mathcal{X}$, $\|P_1^{\pi_1}(\cdot|x) - P_2^{\pi_2}(\cdot|x)\|_1 \leq g_x$, we have*

$$(\mathbf{I} - \gamma P_1^{\pi_1})^{-1} \left(r_1^{\pi_1} - r_2^{\pi_2} - \frac{\gamma R_{\max}}{1 - \gamma} g \right) \leq V_{\mathcal{M}_1}^{\pi_1} - V_{\mathcal{M}_2}^{\pi_2} \leq (\mathbf{I} - \gamma P_1^{\pi_1})^{-1} \left(r_1^{\pi_1} - r_2^{\pi_2} + \frac{\gamma R_{\max}}{1 - \gamma} g \right),$$

where g is the vector of g_x 's. Moreover, the above inequalities are tight.

Proof. The difference between the two value functions can be written as follows:

$$\begin{aligned} V_{\mathcal{M}_1}^{\pi_1} - V_{\mathcal{M}_2}^{\pi_2} &= r_1^{\pi_1} + \gamma P_1^{\pi_1} V_{\mathcal{M}_1}^{\pi_1} - r_2^{\pi_2} - \gamma P_2^{\pi_2} V_{\mathcal{M}_2}^{\pi_2} \\ &= r_1^{\pi_1} + \gamma P_1^{\pi_1} V_{\mathcal{M}_1}^{\pi_1} - r_2^{\pi_2} - \gamma P_2^{\pi_2} V_{\mathcal{M}_2}^{\pi_2} + \gamma P_1^{\pi_1} V_{\mathcal{M}_2}^{\pi_2} - \gamma P_1^{\pi_1} V_{\mathcal{M}_2}^{\pi_2} \\ &= (r_1^{\pi_1} - r_2^{\pi_2}) + \gamma P_1^{\pi_1} (V_{\mathcal{M}_1}^{\pi_1} - V_{\mathcal{M}_2}^{\pi_2}) + \gamma (P_1^{\pi_1} - P_2^{\pi_2}) V_{\mathcal{M}_2}^{\pi_2} \\ &= (\mathbf{I} - \gamma P_1^{\pi_1})^{-1} \left[r_1^{\pi_1} - r_2^{\pi_2} + \gamma (P_1^{\pi_1} - P_2^{\pi_2}) V_{\mathcal{M}_2}^{\pi_2} \right]. \end{aligned}$$

Now using the Holder's inequality, for any $x \in \mathcal{X}$, we have

$$|(P_1^{\pi_1}(\cdot|x) - P_2^{\pi_2}(\cdot|x))^\top V_{\mathcal{M}_2}^{\pi_2}| \leq \|P_1^{\pi_1}(\cdot|x) - P_2^{\pi_2}(\cdot|x)\|_1 \|V_{\mathcal{M}_2}^{\pi_2}\|_\infty \leq g_x \|V_{\mathcal{M}_2}^{\pi_2}\|_\infty \leq g_x \frac{R_{\max}}{1 - \gamma}.$$

The proof follows by uniformly bounding $(P_1^{\pi_1} - P_2^{\pi_2})V_{\mathcal{M}_2}^{\pi_2}$ from the above inequality and from the monotonicity of $(\mathbf{I} - \gamma P_1^{\pi_1})^{-1}$. \square

A.3 Proof of Theorem 3

Proof. From Lemma 2 with $\pi_1 = \pi_2 = \pi_S$, $\mathcal{M}_1 = \mathcal{M}^*$, and $\mathcal{M}_2 = \widehat{\mathcal{M}}$, we have

$$\rho(\pi_S, \widehat{\mathcal{M}}) - \frac{\gamma R_{\max}}{1 - \gamma} p_0^\top (\mathbf{I} - \gamma P_{\pi_S}^*)^{-1} e_{\pi_S} \leq \rho(\pi_S, \mathcal{M}^*).$$

Thus, we may write

$$\begin{aligned} \Phi(\pi_S) &\triangleq \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_S, \mathcal{M}^*) \leq \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_S, \widehat{\mathcal{M}}) + \frac{\gamma R_{\max}}{1 - \gamma} p_0^\top (\mathbf{I} - \gamma P_{\pi_S}^*)^{-1} e_{\pi_S} \\ &\stackrel{(a)}{\leq} \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_{\mathcal{M}^*}^*, \widehat{\mathcal{M}}) + \frac{\gamma R_{\max}}{1 - \gamma} p_0^\top (\mathbf{I} - \gamma P_{\pi_S}^*)^{-1} e_{\pi_S} \\ &\stackrel{(b)}{\leq} \frac{\gamma R_{\max}}{1 - \gamma} p_0^\top \left[(\mathbf{I} - \gamma P_{\pi_{\mathcal{M}^*}^*}^*)^{-1} e_{\pi_{\mathcal{M}^*}^*} + (\mathbf{I} - \gamma P_{\pi_S}^*)^{-1} e_{\pi_S} \right] \\ &\stackrel{(c)}{\leq} \frac{2\gamma R_{\max}}{(1 - \gamma)^2} \|e\|_\infty. \end{aligned}$$

(a) Comes from the optimality of π_S in $\widehat{\mathcal{M}}$.

(b) This is the application of Lemma 2 with policy $\pi_1 = \pi_2 = \pi_{\mathcal{M}^*}^*$, $\mathcal{M}_1 = \mathcal{M}^*$, and $\mathcal{M}_2 = \widehat{\mathcal{M}}$.

(c) This is from the fact that for any policy π , we have $\|p_0^\top (\mathbf{I} - \gamma P_\pi^*)^{-1}\|_1 = 1/(1 - \gamma)$, and from the application of the Holder's inequality. \square

A.4 Proof of Theorem 4

Proof. To prove the safety of π_{Ra} and bound its performance loss, we need to upper and lower bound the difference between the performance of any policy π in the true MDP \mathcal{M}^* and its performance in $\widetilde{\mathcal{M}}$, i.e., $\rho(\pi, \mathcal{M}^*) - \rho(\pi, \widetilde{\mathcal{M}})$. These upper and lower bounds are obtained by applying Lemma 10 with $\pi_1 = \pi_2 = \pi$, $\mathcal{M}_1 = \mathcal{M}^*$, and $\mathcal{M}_2 = \widetilde{\mathcal{M}}$ as follows

$$\rho(\pi, \mathcal{M}^*) - \rho(\pi, \widetilde{\mathcal{M}}) \geq p_0^\top (\mathbf{I} - \gamma P_\pi^*)^{-1} \left(r^\pi - \widehat{r}^\pi - \frac{\gamma R_{\max}}{1 - \gamma} e_\pi \right) \geq 0, \quad (9)$$

where the second inequality in (9) follows from the definition of the adjusted reward function \widehat{r} , and the fact that $(\mathbf{I} - \gamma P_\pi^*)^{-1}$ is monotone and p_0 is non-negative. Similarly, the upper-bound is

$$\rho(\pi, \mathcal{M}^*) - \rho(\pi, \widetilde{\mathcal{M}}) \leq \frac{2\gamma R_{\max}}{1 - \gamma} p_0^\top (\mathbf{I} - \gamma P_\pi^*)^{-1} e_\pi = \frac{2\gamma R_{\max}}{(1 - \gamma)^2} \|e_\pi\|_{1, u_\pi^{\mathcal{M}^*}}, \quad (10)$$

where $u_\pi^{\mathcal{M}^*} = (1 - \gamma)p_0^\top (\mathbf{I} - \gamma P_\pi^*)^{-1}$ is the normalized state occupancy frequency of policy π in the true MDP \mathcal{M}^* .

To prove the safety of the returned policy π_{Ra} , consider the two cases on Line 4 of Algorithm 1. When the condition is satisfied, we have $\rho(\pi_B, \mathcal{M}^*) < \rho(\pi_0, \widetilde{\mathcal{M}}) \leq \rho(\pi_0, \mathcal{M}^*)$, where the second inequality comes from (9), and thus, the policy $\pi_{Ra} = \pi_0$ is safe. When the condition is violated, then π_{Ra} is simply π_B , which is safe by definition.

To derive a bound on the performance loss of the returned policy π_{Ra} , consider also the two cases on Line 4 of Algorithm 1. When the condition is satisfied, using (9), we have

$$\Phi(\pi_{Ra}) = \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_0, \mathcal{M}^*) \leq \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_0, \widetilde{\mathcal{M}}),$$

and when the condition is violated, we have

$$\Phi(\pi_{Ra}) = \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_B, \mathcal{M}^*).$$

Since the condition is satisfied on Line 4 of Algorithm 1 when $\rho(\pi_0, \widetilde{\mathcal{M}}) > \rho(\pi_B, \mathcal{M}^*)$, we may write

$$\Phi(\pi_{Ra}) \leq \min \left\{ \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_0, \widetilde{\mathcal{M}}), \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_B, \mathcal{M}^*) \right\}.$$

Note that we may write the following inequalities for the first term in the minimum

$$\rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_0, \widetilde{\mathcal{M}}) \stackrel{(a)}{\leq} \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_{\mathcal{M}^*}^*, \widetilde{\mathcal{M}}) \stackrel{(b)}{\leq} \frac{2\gamma R_{\max}}{(1-\gamma)^2} \|e_{\pi_{\mathcal{M}^*}^*}\|_{1, u_{\mathcal{M}^*}^*},$$

where **(a)** follows from π_0 being an optimal policy of MDP $\widetilde{\mathcal{M}}$, **(b)** is from (10) with $\pi = \pi_{\mathcal{M}^*}^*$, and $u_{\mathcal{M}^*}^*$ is the normalized state occupancy frequency of the optimal policy $\pi_{\mathcal{M}^*}^*$. This proves the theorem. \square

A.5 Proof of Theorem 5

Proof. To prove the safety of π_R and bound its performance loss, we need to upper and lower bound the difference between the performance of any policy π in the true MDP \mathcal{M}^* and its worst-case performance, $\min_{P \in \mathcal{U}(\widehat{P}, e)} \rho(\pi, \mathcal{M}(P))$. Since $P^* \in \mathcal{U}(\widehat{P}, e)$ from Assumption 1, we have

$$\min_{P \in \mathcal{U}(\widehat{P}, e)} \rho(\pi, \mathcal{M}(P)) \leq \rho(\pi, \mathcal{M}^*). \quad (11)$$

Now let \bar{P} be the minimizer in $\min_{P \in \mathcal{U}(\widehat{P}, e)} \rho(\pi, \mathcal{M}(P))$. The minimizer exists because of the continuity and compactness of the uncertainty set. From Assumption 1 and the construction of $\mathcal{U}(\widehat{P}, e)$, for any $(x, a) \in \mathcal{X} \times \mathcal{A}$, we have

$$\begin{aligned} \|\bar{P}(\cdot|x, a) - P^*(\cdot|x, a)\|_1 &\leq \|\bar{P}(\cdot|x, a) - \widehat{P}(\cdot|x, a) + \widehat{P}(\cdot|x, a) + P^*(\cdot|x, a)\|_1 \\ &\leq \|\bar{P}(\cdot|x, a) - \widehat{P}(\cdot|x, a)\|_1 + \|\widehat{P}(\cdot|x, a) + P^*(\cdot|x, a)\|_1 \\ &\leq 2e(x, a). \end{aligned}$$

Applying Lemma 10 with $\pi_1 = \pi_2 = \pi$, $\mathcal{M}_1 = \mathcal{M}^*$, and $\mathcal{M}_2 = \mathcal{M}(\bar{P})$, we obtain

$$\rho(\pi, \mathcal{M}^*) - \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)) \leq \frac{2\gamma R_{\max}}{1-\gamma} p_0^\top (\mathbf{I} - \gamma P_\pi^*)^{-1} e_\pi = \frac{2\gamma R_{\max}}{(1-\gamma)^2} \|e_\pi\|_{1, u_{\mathcal{M}^*}^\pi}, \quad (12)$$

where $\pi_{\mathcal{M}^*}^\pi = (1-\gamma)p_0^\top (\mathbf{I} - \gamma P_\pi^*)^{-1}$ is the normalized state occupancy frequency of policy π in the true MDP \mathcal{M}^* .

To prove the safety of the returned policy π_R , consider the two cases on Line 4 of Algorithm 2. When the condition is satisfied, we have $\rho(\pi_B, \mathcal{M}^*) < \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi_0, \mathcal{M}(P)) \leq \rho(\pi_0, \mathcal{M}^*)$, where the second inequality comes from (11), and thus, the policy $\pi_R = \pi_0$ is safe. When the condition is violated, then π_R is simply π_B , which is safe by definition.

To derive a bound on the performance loss of the returned policy π_R , consider also the two cases on Line 4 of Algorithm 2. When the condition is satisfied, using (11), we have

$$\Phi(\pi_R) = \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_0, \mathcal{M}^*) \leq \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi_0, \mathcal{M}(P)),$$

and when the condition is violated, we have

$$\Phi(\pi_R) = \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_B, \mathcal{M}^*).$$

Since the condition is satisfied on Line 4 of Algorithm 2 when $\min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi_0, \mathcal{M}(P)) > \rho(\pi_B, \mathcal{M}^*)$, we may write

$$\Phi(\pi_R) \leq \min \left\{ \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi_0, \mathcal{M}(P)) \quad \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_B, \mathcal{M}^*) \right\}.$$

Note that we may write the following inequalities for the first term in the minimum

$$\begin{aligned} \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi_0, \mathcal{M}(P)) &\stackrel{(a)}{\leq} \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}(P)) \\ &\stackrel{(b)}{\leq} \frac{2\gamma R_{\max}}{(1-\gamma)^2} \|e_{\pi_{\mathcal{M}^*}^*}\|_{1, u_{\mathcal{M}^*}^*}, \end{aligned}$$

where (a) follows from π_0 being the maximizer in solving the robust MDP, (b) is from (12) with $\pi = \pi_{\mathcal{M}^*}^*$, and $u_{\mathcal{M}^*}^*$ is the normalized state occupancy frequency of the optimal policy $\pi_{\mathcal{M}^*}^*$. \square

A.6 Proof of Proposition 6

Proof. Let \bar{P} be the minimizer of $\min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P))$. Then from Lemma 10 for each π :

$$\rho(\pi, \mathcal{M}(\bar{P})) \geq \rho(\pi, \widehat{\mathcal{M}}) - \frac{\gamma R_{\max}}{1-\gamma} p_0^\top (\mathbf{I} - \gamma \widehat{P}_\pi)^{-1} e_\pi = \rho(\pi, \widetilde{\mathcal{M}}),$$

The last inequality holds because $\widetilde{\mathcal{M}}$ differs from $\widehat{\mathcal{M}}$ only in its reward function $\hat{r}^\pi = r^\pi - \frac{\gamma R_{\max}}{1-\gamma} e_\pi$ (see (1)). \square

A.7 The Saddle Point Solution Algorithm

Now we turn to the solution algorithm of (6). Notice that for every fixed $\lambda \geq 0$, the inner optimization problem in (6) is a robust MDP problem with (s, a) -rectangular uncertainty. We now summarize the standard results of robust value iteration from [16, 9]. For any $(x, y) \in \mathcal{X} \times \mathcal{X}$, define the λ -parametrized robust Bellman's operator $\mathcal{T}_\lambda : \mathbb{R}^{|\mathcal{X} \times \mathcal{X}|} \rightarrow \mathbb{R}^{|\mathcal{X} \times \mathcal{X}|}$ as follows:

$$\mathcal{T}_\lambda[V](x, y) = \max_{a \in \mathcal{A}} \left\{ r_\lambda^C(x, y, a) + \gamma \min_{P \in \mathcal{U}(\hat{P}, e)} \sum_{(x', y') \in \mathcal{X} \times \mathcal{X}} \mathcal{P}^C(x', y' | x, y, a) V(x', y') \right\}.$$

Since the robust Bellman operator is a contraction mapping, its unique fixed point solution equals to the optimal robust Lagrangian function, i.e., $\mathcal{V}_\lambda(x, y) = \mathcal{T}_\lambda[\mathcal{V}_\lambda](x, y)$, $\forall (x, y) \in \mathcal{X} \times \mathcal{X}$ and

$$\sum_{x, x \in \mathcal{X}} p_0(x) p_0(x) \mathcal{V}_\lambda(x, y) = \min_{P \in \mathcal{U}(\hat{P}, e)} \max_{\pi \in \Pi_S^C} L(\mathcal{P}^C, \pi, \lambda). \quad (13)$$

Thus the robust value function $\mathcal{V}_\lambda(x, y)$ can be calculated by robust value iteration [16],

$$\mathcal{V}_{\lambda, 0}(x, y) = \mathcal{V}_0(x, y), \quad \mathcal{V}_{\lambda, N+1}(x, y) = \mathcal{T}_\lambda[\mathcal{V}_{\lambda, N}](x, y), \quad \forall (x, y) \in \mathcal{X} \times \mathcal{X}, \quad N \in \{0, 1, 2, \dots\},$$

where the λ -parametrized optimal policy $\pi_\lambda^* : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{A}$ and worst case transition probability P^* form a *minimax* saddle point of the fixed point solution for $\lambda \geq 0$ ². Furthermore, the following sub-gradient descent algorithm finds the Lagrange multiplier of problem (6):

1. Find the following Lagrange multiplier update:

$$\lambda^{(j+1)} = \left(\lambda^{(j)} - \alpha^{(j)} (\rho(\pi_{\lambda^{(j)}}^*, \mathcal{M}(P^*)) - \rho(\pi_B, \mathcal{M}^*)) \right)^+$$

where the step length $\alpha^{(j)}$ is non-summable, square-summable.³

2. Define $f_{\min}^{(j+1)} = \min(f_{\min}^{(j)}, f(\lambda^{(j+1)}))$ as the best dual function estimate where $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is the robust Lagrangian function at $\pi = \nu_\lambda$, i.e.,

$$f(\lambda) = \max_{\pi \in \Pi_S^C} \min_{P \in \mathcal{U}(\hat{P}, e)} L(\pi, \lambda).$$

Update the Lagrange multiplier estimate as follows:

$$\lambda^{(j+1)} \leftarrow \begin{cases} \lambda^{(j+1)} & \text{if } f_{\min}^{(j+1)} = f(\lambda^{(j+1)}) \\ \lambda^{(j)} & \text{otherwise} \end{cases}.$$

²Note that λ is a linear function of $\min_{P \in \mathcal{U}(\hat{P}, e)} L(\pi, \lambda)$ and the worst case minimization is only taken with respect to the constraint component. Then P^* is independent of λ .

³The step-size condition satisfies $\alpha^{(j)} \geq 0$, $\sum_{j=0}^{\infty} \alpha^{(j)} = \infty$ and $\sum_{j=0}^{\infty} (\alpha^{(j)})^2 < \infty$.

The following Lemma shows that the solution of the above procedure converges to the solution of $\min_{\lambda \geq 0} f(\lambda)$.

Lemma 11. *Let λ^* be the minimizer of $\min_{\lambda \geq 0} f(\lambda)$. The solution of the projected sub-gradient descent algorithm converges to λ^* , i.e., $\lim_{j \rightarrow \infty} \lambda^{(j)} = \lambda^*$.*

Proof. Since $L(\pi, \lambda)$ is a linear function of λ for any given \mathcal{P}^C and π , we have that $f(\lambda) = \max_{\pi \in \Pi_S^C} \min_{P \in \mathcal{U}(\hat{P}, e)} L(\pi, \lambda)$ is a convex function in λ . By the envelop theorem of mathematical economics [15], for $\lambda \geq 0$ one can write

$$\frac{df(\lambda)}{d\lambda} = \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)) - \rho(\pi_B, \mathcal{M}^*).$$

Now, we show that the proposed sub-gradient descent algorithm converges to the optimizer of $\min_{\lambda \geq 0} f(\lambda)$. First, one can write the λ iterate as $\lambda^{(j+1)} = (\bar{\lambda}^{(j+1)})^+$ where

$$\bar{\lambda}^{(j+1)} := \lambda^{(j)} - \alpha^{(j)} \frac{df(\lambda)}{d\lambda} \Big|_{\lambda=\lambda^{(j)}}.$$

Since the projection operator is non-expansive, one obtains $(\lambda^{(j+1)} - \lambda^*)^2 \leq (\bar{\lambda}^{(j+1)} - \lambda^*)^2$. Furthermore, the following expression holds:

$$\begin{aligned} (\lambda^{(j+1)} - \lambda^*)^2 &\leq (\bar{\lambda}^{(j+1)} - \lambda^*)^2 \\ &= \left(\lambda^{(j)} - \alpha^{(j)} \frac{df(\lambda)}{d\lambda} \Big|_{\lambda=\lambda^{(j)}} - \lambda^* \right)^2 \\ &= (\lambda^{(j)} - \lambda^*)^2 - 2\alpha^{(j)} (\lambda^{(j)} - \lambda^*) \frac{df(\lambda)}{d\lambda} \Big|_{\lambda=\lambda^{(j)}} + \left(\alpha^{(j)} \right)^2 \left(\frac{df(\lambda)}{d\lambda} \Big|_{\lambda=\lambda^{(j)}} \right)^2 \\ &\leq (\lambda^{(j)} - \lambda^*)^2 - 2\alpha^{(j)} (f(\lambda^{(j)}) - f(\lambda^*)) + \left(\alpha^{(j)} \right)^2 \left(\frac{df(\lambda)}{d\lambda} \Big|_{\lambda=\lambda^{(j)}} \right)^2, \end{aligned}$$

The inequality is due to the fact that for any $\lambda \geq 0$, convexity of $f(\lambda)$ implies that

$$f(\lambda) - f(\lambda^*) \geq (\lambda - \lambda^*) \frac{df(\lambda)}{d\lambda}.$$

This further implies

$$(\lambda^{(j+1)} - \lambda^*)^2 \leq (\lambda^{(0)} - \lambda^*)^2 - \sum_{q=0}^j 2\alpha^{(q)} (f(\lambda^{(q)}) - f(\lambda^*)) + \left(\alpha^{(q)} \right)^2 \left(\frac{df(\lambda)}{d\lambda} \Big|_{\lambda=\lambda^{(q)}} \right)^2.$$

Since $(\lambda^{(j+1)} - \lambda^*)^2$ is a positive quantity and $(\lambda^{(0)} - \lambda^*)^2$ is bounded, this further implies

$$2 \sum_{q=0}^j \alpha^{(q)} (f(\lambda^{(q)}) - f(\lambda^*)) \leq (\bar{\lambda}^{(0)} - \lambda^*)^2 + \sum_{q=0}^j \left(\alpha^{(q)} \right)^2 \left(\frac{df(\lambda)}{d\lambda} \Big|_{\lambda=\lambda^{(q)}} \right)^2.$$

By defining $f_{\min}^{(j)} = \min_{q \in \{0, \dots, j\}} f(\lambda^{(q)})$, the above expression implies

$$f_{\min}^{(j)} - f(\lambda^*) \leq \frac{1}{\sum_{q=0}^j \alpha^{(q)}} \left((\lambda^{(0)} - \lambda^*)^2 + \sum_{q=0}^j \left(\alpha^{(q)} \right)^2 \left(\frac{df(\lambda)}{d\lambda} \Big|_{\lambda=\lambda^{(q)}} \right)^2 \right).$$

The step-size rule of $\alpha^{(q)}$ ensures that the numerator is bounded and the denominator goes to infinity as $j \rightarrow \infty$. This implies that for any $\epsilon > 0$, there exists a constant $N(\epsilon)$ such that for any $j > N(\epsilon)$, $f(\lambda^*) \leq f_{\min}^{(j)} \leq f(\lambda^*) + \epsilon$. In other words, the sequence $\lambda^{(j)}$ converges to the global minimum λ^* of $f(\lambda)$. \square

Combining all previous arguments, the saddle point solution of (6) is given by $(\pi_\lambda^*, \lambda^*)$.

A.8 Proof of Theorem 7

Proof. Given Assumption 1, if the solution $\min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda)$ is lower-bounded, then weak duality implies that the primal Lagrangian $\max_{\pi \in \Pi_S^A} \min_{\lambda \geq 0} L(\pi, \lambda)$ is also lower-bounded, which further implies that π_{AR} is a *safe* policy. Otherwise, Algorithm 3 returns the baseline policy π_B . This concludes that the policy π_{AR} returned by Algorithm 3 is safe.

For the performance loss bound, without loss of generality it is analyzed based on MDP $\mathcal{M}_{\lambda_1, \lambda_2}^A(P)$ with $\lambda_1 = 1$ and $\lambda_2 = 0$, where $\mathcal{M}_{1,0}^A(P) = \mathcal{M}(P)$ and $\mathcal{M}_{1,0}^A(P^*) = \mathcal{M}^*$. The proof follows with arguments identical to those in the proof of Theorem 5 and is omitted for the sake of brevity. \square

A.9 Proof of Proposition 8

Proof. Suppose that Algorithm 2 returns a safe policy other than π_B , i.e., $\pi_R \neq \pi_B$. Since the class of policies Π_S^A is dominating for the optimization problem (2), we may write (2) as

$$\max_{\pi \in \Pi_S^{AR}} \rho(\pi, \widehat{\mathcal{M}}), \quad (14)$$

where the feasible policy set Π_S^{AR} is defined as

$$\Pi_S^{AR} := \left\{ \pi \in \Pi_S^A : \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)) \geq \rho(\pi_B, \mathcal{M}^*) \right\}.$$

This leads to the primal Lagrangian formulation $\max_{\pi \in \Pi_S^A} \min_{\lambda \geq 0} \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda)$.

Since $\pi_R \neq \pi_B$ is a safe policy and the feasible set Π_S^{AR} is non-empty, the solution of the primal Lagrangian is equal to the solution of (14). Furthermore by weak duality, we have

$$\max_{\pi \in \Pi_S^A} \min_{\lambda \geq 0} \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda) \leq \underbrace{\min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda)}_{\text{solution of Algorithm 3}}.$$

On the other hand, consider the following optimization problem:

$$\max_{\pi \in \Pi_S^{AR}} \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)). \quad (15)$$

When $\pi_R \neq \pi_B$ is a safe policy, the feasible set Π_S^{AR} is non-empty, and we may write the Lagrangian function of (15) as

$$\mathcal{L}(\pi, \lambda) = \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)) + \lambda \left(\min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)) - \rho(\pi_B, \mathcal{M}^*) \right)$$

Since the feasible set is non-empty, we may drop the constraint in (15) and this problem becomes equivalent to

$$\max_{\pi \in \Pi_S^A} \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)),$$

where the dominating class of policies of this problem is stationary Markovian Π_S . Thus the primal Lagrangian duality formulation implies that

$$\max_{\pi \in \Pi_S^A} \min_{\lambda \geq 0} \mathcal{L}(\pi, \lambda) = \max_{\pi \in \Pi_S^A} \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)) = \max_{\pi \in \Pi_S} \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)),$$

i.e., $\max_{\pi \in \Pi_S^A} \min_{\lambda \geq 0} \mathcal{L}(\pi, \lambda)$ equals to the solution of Algorithm 2. Since the objective function of (15) is lower than that of (14) and both problems share the same feasible set, it is obvious that

$$\max_{\pi \in \Pi_S^A} \min_{\lambda \geq 0} \mathcal{L}(\pi, \lambda) \leq \max_{\pi \in \Pi_S^A} \min_{\lambda \geq 0} \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda).$$

Combining these arguments, we have just showed that if π_R is a safe policy, then

$$\rho(\pi_B, \mathcal{M}^*) \leq \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi_R, \mathcal{M}(P)) \leq \min_{\lambda \geq 0} \max_{\pi \in \Pi_S^A} \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda) \leq \min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi_{AR}, \lambda^*),$$

where (π_{AR}, λ^*) is the *maximin* saddle-point solution of $\min_{P \in \mathcal{U}(\hat{P}, e)} L_P(\pi, \lambda)$. \square

A.10 Proof of Theorem 9

Proof. Consider an arbitrary $\tilde{P} \in \mathcal{U}(\hat{P}, e)$. Then from Assumption 1 and the construction of $\mathcal{U}(\hat{P}, e)$, we have:

$$\begin{aligned} \|\tilde{P}(\cdot|x, a) - P^*(\cdot|x, a)\|_1 &\leq \|\tilde{P}(\cdot|x, a) - \hat{P}(\cdot|x, a) + \hat{P}(\cdot|x, a) + P^*(\cdot|x, a)\|_1 \\ &\leq \|\tilde{P}(\cdot|x, a) - \hat{P}(\cdot|x, a)\|_1 + \|\hat{P}(\cdot|x, a) + P^*(\cdot|x, a)\|_1 \\ &\leq 2e(x, a). \end{aligned}$$

Then, using Lemma 10 with the above difference between \tilde{P} and P^* , we get for any policy π that:

$$\max_{P \in \mathcal{U}(\hat{P}, e)} \left| \rho(\pi, P) - \rho(\pi, P^*) \right| \leq \frac{2\gamma R_{\max}}{1 - \gamma} p_0^\top (\mathbf{I} - P_\pi^*) e_\pi = \frac{2\gamma R_{\max}}{(1 - \gamma)^2} \|e_\pi\|_{1, u}, \quad (16)$$

where u_π is the normalized state occupancy frequency for policy π defined as:

$$u = (1 - \gamma)(\mathbf{I} - \gamma P_\pi^\top)^{-1} p_0.$$

To prove the safety of π_I , note that the objective in (7) is always non-negative since π_B is feasible. Then we get the safety condition by simple algebraic manipulation as follows:

$$\begin{aligned} \min_{P \in \mathcal{U}(\hat{P}, e)} \left(\rho(\pi_I, P) - \rho(\pi_B, P) \right) &\geq 0 \\ \rho(\pi_I, P^*) &\geq \rho(\pi_B, P^*) \end{aligned}$$

The safety of the policy π_I also implies that its performance loss is bounded by the performance loss of the base policy:

$$\rho(\pi_I, P^*) - \rho(\pi_I, P^*) \leq \rho(\pi_I, P^*) - \rho(\pi_B, P^*) \quad (17)$$

Now we ready to show a bound on the performance loss of π_I by lower bounding $\rho(\pi_I, P^*)$ as follows:

$$\rho(\pi_I, P^*) = \rho(\pi_I, P^*) - \rho(\pi_B, P^*) + \rho(\pi_B, P^*) \geq \min_{P \in \mathcal{U}(\hat{P}, e)} \left(\rho(\pi_I, P) - \rho(\pi_B, P) \right) + \rho(\pi_B, P^*).$$

From the optimality of π_I , we further get:

$$\begin{aligned} \min_{P \in \mathcal{U}(\hat{P}, e)} \left(\rho(\pi_I, P) - \rho(\pi_B, P) \right) &\geq \min_{P \in \mathcal{U}(\hat{P}, e)} \left(\rho(\pi^*, P) - \rho(\pi_B, P) \right) \\ &\geq \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi^*, P) - \max_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi_B, P). \end{aligned}$$

Putting the above together and some simple algebraic manipulation subtracting and adding $\rho(\pi^*, P^*)$, we get:

$$\rho(\pi^*, P^*) - \rho(\pi_I, P^*) \leq \max_{P \in \mathcal{U}(\hat{P}, e)} \left(\rho(\pi^*, P^*) - \rho(\pi^*, P) \right) + \max_{P \in \mathcal{U}(\hat{P}, e)} \left(\rho(\pi_B, P) - \rho(\pi_B, P^*) \right).$$

The bound in the theorem then follows by bounding the maximization terms above using (16) and combining the above inequality with (17).

Figure 4 depicts an example demonstrating the tightness of the bound. The initial state is s_0 , actions are a_1, a_2 , and the transitions are deterministic. ξ^* denotes the true transitions, ξ_1 denotes the worst case in $\mathcal{U}(\hat{P}, e)$, and the leafs shows the returns of the remainder of the MDP and are assumed to be known with certainty. The value for ϵ is given by (16). That is $\hat{\xi}$ represents the sampled transition probability, it would be halfway between ξ^* and ξ_1 . \square




Figure 4: Example showing tightness of bound in Theorem 9.

B Alternative Bounds

B.1 An Alternative Bound on the Performance Loss of π_{Ra}

Corollary 12. *Given Assumption 1, the performance loss $\Phi(\pi_{Ra})$ also satisfies*

$$\Phi(\pi_{Ra}) \leq \min \left\{ \frac{BR(\pi_{Ra})}{1 - \gamma} + \max_{\pi \in \Pi_S} \frac{2\gamma R_{\max}}{(1 - \gamma)^2} \|e_{\pi, \tilde{\mathcal{M}}}\|_{1, \pi_{\tilde{\mathcal{M}}}}, \Phi(\pi_B) \right\},$$

where $BR(\pi_{Ra}) = \max_{x \in \mathcal{X}} |\hat{T}[\hat{V}^{\pi_{Ra}}](x) - \hat{V}^{\pi_{Ra}}(x)|$ is the Bellman residual w.r.t. Bellman operator $\hat{T}[V](x) = \max_{a \in \mathcal{A}} \left\{ \hat{r}(x, a) + \gamma \sum_{x' \in \mathcal{X}} \hat{P}(x'|x, a) V(x') \right\}$, value function $\hat{V}^{\pi_{Ra}}(x) = \rho(\pi_{Ra}, \tilde{\mathcal{M}})$ at $x_0 = x$ and $\pi_{\tilde{\mathcal{M}}}^{\pi}$ is the normalized state occupancy frequency of the policy π .

Proof. Similar to the proof of Theorem 12, we notice that

$$\rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \rho(\pi_0, \tilde{\mathcal{M}}) = \underbrace{\rho(\pi_{\mathcal{M}^*}^*, \mathcal{M}^*) - \max_{\pi \in \Pi_S} \rho(\pi, \tilde{\mathcal{M}})}_{(a)} + \underbrace{\max_{\pi \in \Pi_S} \rho(\pi, \tilde{\mathcal{M}}) - \rho(\pi_0, \tilde{\mathcal{M}})}_{(b)} \quad (18)$$

First we prove an upper bound for (a). Recall the Bellman operator \hat{T} for value function $V : \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}^{\mathcal{X}}$ as $\hat{T}[V](y) = \max_{a \in \mathcal{A}} \left\{ \hat{r}(y, a) + \gamma \sum_{y' \in \mathcal{X}} \hat{P}(y'|y, a) V(y') \right\}$. Also define the value function $\hat{V}^{\pi_0}(x)$ as $\rho(\pi_0, \tilde{\mathcal{M}})$ and the optimal value function $\hat{V}(x)$ as $\max_{\pi \in \Pi_S} \rho(\pi, \tilde{\mathcal{M}})$, when the initial state is $x_0 = x$. By applying the contraction mapping property on $\hat{T}[\hat{V}^{\pi_0}](y) - \hat{V}^{\pi_0}(y)$ for any $y \in \mathcal{X}$ and combining with the definition of Bellman residual, one obtains

$$\hat{T}^2[\hat{V}^{\pi_0}](y) - \hat{T}[\hat{V}^{\pi_0}](y) \leq \gamma BR(\pi_0).$$

By an induction argument, the above expression becomes

$$\hat{T}^N[\hat{V}^{\pi_0}](y) - \hat{T}^{N-1}[\hat{V}^{\pi_0}](y) \leq \gamma^{N-1} BR(\pi_0),$$

for which by a telescoping sum, it further implies

$$\widehat{T}^N[\widehat{V}^{\pi_0}](y) - \widehat{V}^{\pi_0}(y) = \sum_{k=1}^N \widehat{T}^k[\widehat{V}^{\pi_0}](y) - \widehat{T}^{k-1}[\widehat{V}^{\pi_0}](y) \leq \sum_{k=1}^N \gamma^{k-1} \text{BR}(\pi_0).$$

By letting $N \rightarrow \infty$ and noticing that $\lim_{N \rightarrow \infty} \widehat{T}^N[\widehat{V}^{\pi_0}](y) = \widehat{V}(y)$, one finally obtains

$$(b) \leq \left| \max_{\pi \in \Pi_S} \rho(\pi, \widetilde{\mathcal{M}}) - \rho(\pi_0, \widetilde{\mathcal{M}}) \right| = \left| \sum_{y \in \mathcal{X}} P_0(y) \left(\widehat{V}(y) - \widehat{V}^{\pi_0}(y) \right) \right| \leq \lim_{N \rightarrow \infty} \sum_{k=1}^N \gamma^{k-1} \text{BR}(\pi_0) = \frac{\text{BR}(\pi_0)}{1 - \gamma}. \quad (19)$$

For an upper bound in (a), by interchanging \mathcal{M} and $\widetilde{\mathcal{M}}$ in the derivation of 10 and taking maximum on both sides, we also have,

$$(a) = \max_{\pi \in \Pi_S} \rho(\pi, \mathcal{M}^*) - \max_{\pi \in \Pi_S} \rho(\pi, \widetilde{\mathcal{M}}) \leq \max_{\pi \in \Pi_S} \frac{2\gamma R_{\max}}{1 - \gamma} p_0^\top (\mathbf{I} - \gamma \widehat{P}_\pi)^{-1} e_\pi \leq \max_{\pi \in \Pi_S} \frac{2\gamma R_{\max}}{(1 - \gamma)^2} \|e_{\pi, \mathcal{M}}\|_{1, \pi_{\mathcal{M}}^\pi}.$$

Then the proof is completed by combining both parts of the above arguments. \square

B.2 An Alternative Bound on the Performance Loss of π_R

Corollary 13. *Given Assumption 1, the performance loss $\Phi(\pi_R)$ also satisfies*

$$\Phi(\pi_R) \leq \min \left\{ \frac{\text{BR}(\pi_R)}{1 - \gamma} + \max_{\pi \in \Pi_S} \min_{P \in \mathcal{U}(\widehat{P}, e)} \frac{2\gamma R_{\max}}{(1 - \gamma)^2} \|e_{\pi, \mathcal{M}(P)}\|_{1, \pi_{\mathcal{M}(P)}}, \Phi(\pi_B) \right\},$$

where $\text{BR}(\pi_R) = \max_{x \in \mathcal{X}} |\mathcal{T}[V^{\pi_R}](x) - V^{\pi_R}(x)|$ is the Bellman residual w.r.t. Bellman operator $\mathcal{T}[V](x) = \max_{a \in \mathcal{A}} \left\{ r(x, a) + \gamma \min_{P \in \mathcal{U}(\widehat{P}, e)} \sum_{x' \in \mathcal{X}} P(x'|x, a) V(x') \right\}$, value function $V^{\pi_R}(x) = \min_{P \in \mathcal{U}(\widehat{P}, e)} \rho(\pi_R, \mathcal{M}(P))$ at $x_0 = x$ and $u_{\mathcal{M}(P)}^\pi$ is the normalized state occupancy frequency of the policy π .

Proof. For the proof of the alternative performance bound, by following the same analysis as in the derivation of (18), replacing the bellman operator \widehat{T} with the robust Bellman operator:

$$\mathcal{T}[V](y) = \max_{a \in \mathcal{A}} \left\{ r(y, a) + \gamma \min_{P \in \mathcal{U}(\widehat{P}, e)} \sum_{y' \in \mathcal{X}} P(y'|y, a) V(y') \right\},$$

and defining the robust value function $V^{\pi_0}(x)$ as $\min_{P \in \mathcal{U}(\widehat{P}, e)} \rho(\pi_0, \mathcal{M}(P))$ when the initial state is $x_0 = x$, we can easily show that

$$\max_{\pi \in \Pi_S} \min_{P \in \mathcal{U}(\widehat{P}, e)} \rho(\pi, \mathcal{M}(P)) - \min_{P \in \mathcal{U}(\widehat{P}, e)} \rho(\pi_0, \mathcal{M}(P)) \leq \frac{\text{BR}(\pi_0)}{1 - \gamma},$$

and

$$\begin{aligned} \max_{\pi \in \Pi_S} \rho(\pi, \mathcal{M}^*) - \max_{\pi \in \Pi_S} \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi, \mathcal{M}(P)) &\leq \max_{\pi \in \Pi_S} \min_{P \in \mathcal{U}(\hat{P}, e)} \frac{2\gamma R_{\max}}{1-\gamma} p_0^\top (\mathbf{I} - \gamma P_\pi)^{-1} e_\pi \\ &\leq \max_{\pi \in \Pi_S} \min_{P \in \mathcal{U}(\hat{P}, e)} \frac{2\gamma R_{\max}}{(1-\gamma)^2} \|e_{\pi_{\mathcal{M}(P)}}\|_{1, u_{\mathcal{M}(P)}^\pi}. \end{aligned}$$

The proof is then completed by combining both of the above results. \square

B.3 An Alternative Bound on the Performance Loss of π_{AR}

Corollary 14. *Given Assumption 1, the performance loss $\Phi(\pi_{AR})$ also satisfies*

$$\Phi(\pi_{AR}) \leq \min \left\{ \frac{BR(\pi_{AR})}{1-\gamma} + \max_{\pi \in \Pi_S} \min_{P \in \mathcal{U}(\hat{P}, e)} \frac{2\gamma R_{\max}}{(1-\gamma)^2} \|e_{\pi_{\mathcal{M}(P)}}\|_{1, \pi_{\mathcal{M}(P)}}, \Phi(\pi_B) \right\}.$$

where $u_{\mathcal{M}(P)}^\pi$ is the normalized state occupancy frequency of the policy π and the Bellman residual for robust MDPs is $BR(\pi_{AR}) = BR_{1,0}(\pi_{AR})$, for which the generic case $(\lambda_1, \lambda_2) \geq 0$, the Bellman residual w.r.t. Bellman operator $\mathcal{T}_{\lambda_1 \lambda_2}[V](x, y) = \max_{a \in \mathcal{A}} \left\{ r_{\lambda_1, \lambda_2}^A(x, y, a) + \gamma \min_{P \in \mathcal{U}(\hat{P}, e)} \sum_{x', y' \in \mathcal{X}} P^A(x', y' | x, y, a) V(x', y') \right\}$ is given by $BR_{\lambda_1, \lambda_2}(\pi_{RS}) = \max_{x, y \in \mathcal{X}} |\mathcal{T}_{\lambda_1 \lambda_2}[V_{\lambda_1 \lambda_2}^{\pi_{AR}}](x, y) - V_{\lambda_1 \lambda_2}^{\pi_{AR}}(x, y)|$ where $V^{\pi_{AR}}(x, y) = \min_{P \in \mathcal{U}(\hat{P}, e)} \rho(\pi_{AR}, \mathcal{M}_{\lambda_1 \lambda_2}^A(P))$ is the value function at $x_0 = x$.

Proof. Since to the proof of Theorem 7, the proof of this corollary follows identical arguments from Corollary 13 and is omitted for the sake of brevity. \square